In this article

Memory Size (RAM)
Memory Bandwidth (RAM)
CPU Core Count
CPU Clock Speed
Workstations With Multiple CPU's
Clusters (Multiple Computers on a Network)
Network Speed and Latency
Cloud Computing
GPU support

TOP

Lumerical hardware specifications guidelines

In this article we intend to provide some general tips and guidelines for users looking to buy new hardware, or who are just interested in understanding how certain computer hardware components may effect the speed of Lumerical simulations on a single workstation and on high speed computers/cluster.

Memory Size (RAM)

The memory size determines the size of the simulation design or project that can run on the computer. It does not affect the simulation speed as long as the entire simulation can fit into the RAM, otherwise the computer may be forced to swap memory to the hard drive (or the application will provide an error). If swapping occurs, the simulation speed will be extremely slow.

Desktop computers nowadays typically have 8-32 GB of RAM, which is sufficient for running a large fraction of simulations. Workstations often have 32-128 GB of RAM, which is sufficient for running almost all simulations. You can **check how much memory is required to run FDTD simulation**.

When running Lumerical simulation, large amounts of data must be continuously transferred between the RAM and CPU. When the memory bus is unable to transfer data fast enough, the processor is forced to wait, limiting the overall speed of your simulation. For example, on a typical desktop computer with 8 cores, the simulation speed might increase by a factor of 2-3x when using eight cores compared to one core. Therefore, when purchasing a computer, having a fast memory bandwidth is very important.

CPU Core Count

Our simulation tools will try and use all of your CPU cores to run as quickly as possible. However, as we learned above, our simulation is memory bandwidth limited, so adding more cores does not always make the simulation go faster. The speedup that most users see when moving to higher core count CPUs is most often due to other (memory related) improvements to the CPU architecture.

Most CPUs' support hyper-threading, which allows the operating system to treat each physical CPU core as two logical cores. This feature does not provide any speed increase for FDTD simulations because the overall performance bottleneck is the data transfer rate between the CPU and RAM, not the actual computing capability of the cores.

CPU Clock Speed

The CPU clock speed is typically not the most important factor for Lumerical simulation speeds. While a faster clock speed does allow each core to run more quickly, the overall simulation speed is limited by the access between CPU and RAM.

Workstations With Multiple CPU's

Workstations with multiple CPU's are a good way to increase the simulation speed. The most important factor is that each CPU has its own memory bus connection to the RAM. As explained above, the data transfer rate between CPU and RAM is the performance bottleneck, so having one memory bus per CPU allows the simulation speed scale very well with the number of CPU's.

> ***Example:***
>
> *The `Intel Xeon Gold 5115` has a '`Maximum memory bandwidth`' of 107 GB/s. Up to 4 of these processors can be installed in a single workstation, which would give*

Optics

References:

- Intel Xeon Gold Processor 65

- Xeon Gold 5115 Processor Specs 63

- Intel Xeon Processors - Value Analysis 65

Clusters (Multiple Computers on a Network)

In applications where a single computer is not enough, multiple computers can be connected over a network to form a "cluster". FDTD Solutions supports 2 modes when running on a cluster: running multiple simulations across a network (**Concurrent Parametric Computing**) and running a single, large complex simulation across multiple computers (**Distributed Computing**).

Network Speed and Latency

- When running a simulation locally on a single computer, the network speed does not have any effect on the simulation speed.
- When running a simulation remotely on a single computer, or **running multiple simulations on a cluster** (ie. a sweep or optimization), network latency has no effect on the simulation speed, and network speed will only affect how fast results can be retrieved.
- When **running a single, large simulation across multiple computers** , the network speed is extremely important - high speed, low latency interconnects such as **InfiniBand** are recommended in such cases.

Cloud Computing

You can use cloud computing services to evaluate the latest hardware before making a purchasing decision, or in cases where a long term hardware investment does not make sense:

- Amazon Web Services: **Instance Types** "General Purpose" and "Compute Optimized". **Running FDTD Solutions on AWS** .
- Microsoft Azure: **Instance Types:** "General Purpose", "Compute Optimized", and "High Performance Computing". **Running FDTD Solutions on Microsoft Azure** .
- **more**...

GPU support

Notes: *All the above examples are not intended to be endorsements of these models or brands. They are simply examples used to illustrate the points described in the page.*

See also

FDTD GPU Solver Information

FDTD Performance Benchmarks

Resource configuration elements and controls

Determining the optimum resources to run FDTD

Related articles

Getting started with running FDTD on GPU

Running FDTD simulations from the design environment

FDTD Performance Benchmarks

System Requirements and Supported Platform

Determining the optimum resources to run FDTD

Optics

Learn

Ansys Innovation Courses
Ansys Learning Hub
Webinars
On-Demand Webinars
White Papers

Get Help

Ansys Learning Forum
Lumerical Ideas Exchange
Ansys Support
How to Register for Support

Evaluate for Free

Lumerical
Zemax
Speos

Hear the Latest from Ansys

Copyright 2025 Ansys Canada Ltd.

Terms of Service Cookie Policy Privacy Policy Legal Notices

Optics